Universal Deep Image Compression via Content-Adaptive Optimization with Adapters: Supplementary Material

Koki Tsubota[1] Hiroaki Akutsu[2] Kiyoharu Aizawa[1]
[1]The University of Tokyo [2]Hitachi, Ltd.
{tsubota,aizawa}@hal.t.u-tokyo.ac.jp, hiroaki.akutsu.cs.@hitachi.com

This supplementary material provides (1) detailed experimental setup of existing adaptive compression methods and (2) whole images of the qualitative results that we are not able to show in the main paper.

1. Detailed Experimental Setup of Existing Adaptive Compression Methods

We describe the detailed experimental setup of four existing adaptive compression methods: (1) Yang *et al*. [3], (2) Lam *et al*. [1], (3) Rozendaal *et al*. [2], and (4) Zou *et al*. [4]. (1) Yang *et al*. performs latent refinement without any adapter, which is equivalent to the first stage of our approach. (2) Lam *et al*. updates the parameters of the biases of the convolutional layers in the decoder after the latent refinement. The parameters are optimized only in terms of distortion. The updated parameters are converted to 64 bits floating point numbers and compressed in the 7z format. The number of updated parameters is 9283. (3) Rozendaal *et al*. updates all the parameters in the decoder and entropy model. The parameters are optimized in terms of rate-distortion. The number of updated parameters is 6.50×10^7. (4) Lam *et al*. uses overfittable multiplicative parameters (OMPs) instead of adapters after the latent refinement. The parameters are optimized only in terms of distortion. The updated parameters are transformed linearly to the range of [0, 255] and quantized to the integer. This was performed to obtain integer values of eight bits and two real values of 32 bits, which are the scale and bias for the linear transformation. The number of updated parameters is 192.

2. Whole Images of Qualitative Results

The comparison results with other compression methods are shown in Figs. 1, 2, 3, and 4. The ablation results on the effectiveness of adapters are shown in Fig. 5.

References

[1] Yat Hong Lam, Alireza Zare, Francesco Cricri, Jani Lainema, and Miska M. Hannuksela. Efficient adaptation of neural network filter for video compression. In *ACMMM*, pages 358–366, Virtual, Oct. 2020.

[2] Ties van Rozendaal, Iris AM Huijben, and Taco Cohen. Overfitting for fun and profit: Instance-adaptive data compression. In *ICLR*, Virtual, May 2021.

[3] Yibo Yang, Robert Bamler, and Stephan Mandt. Improving inference for neural image compression. In *NeurIPS*, Virtual, Dec. 2020.

[4] Nannan Zou, Honglei Zhang, Francesco Cricri, Ramin G. Youvalari, Hamed R. Tavakoli, Jani Lainema, Emre Aksu, Miska Hannuksela, and Esa Rahtu. Adaptation and attention for neural video coding. In *ISM*, pages 240–244, Italy, Nov. 2021.

[5] Renjie Zou, Chunfeng Song, and Zhaoxiang Zhang. The devil is in the details: Window-based attention for image compression. In *CVPR*, USA, June 2022.

Input (Natual image)
PSNR/BPP

Ours
29.6/0.242

JPEG
23.7/0.252

BPG
28.7/0.254

VVC
29.5/0.254

WACNN [5]
29.3/0.247

Figure 1: Qualitative comparison with other compression methods on a natural image.

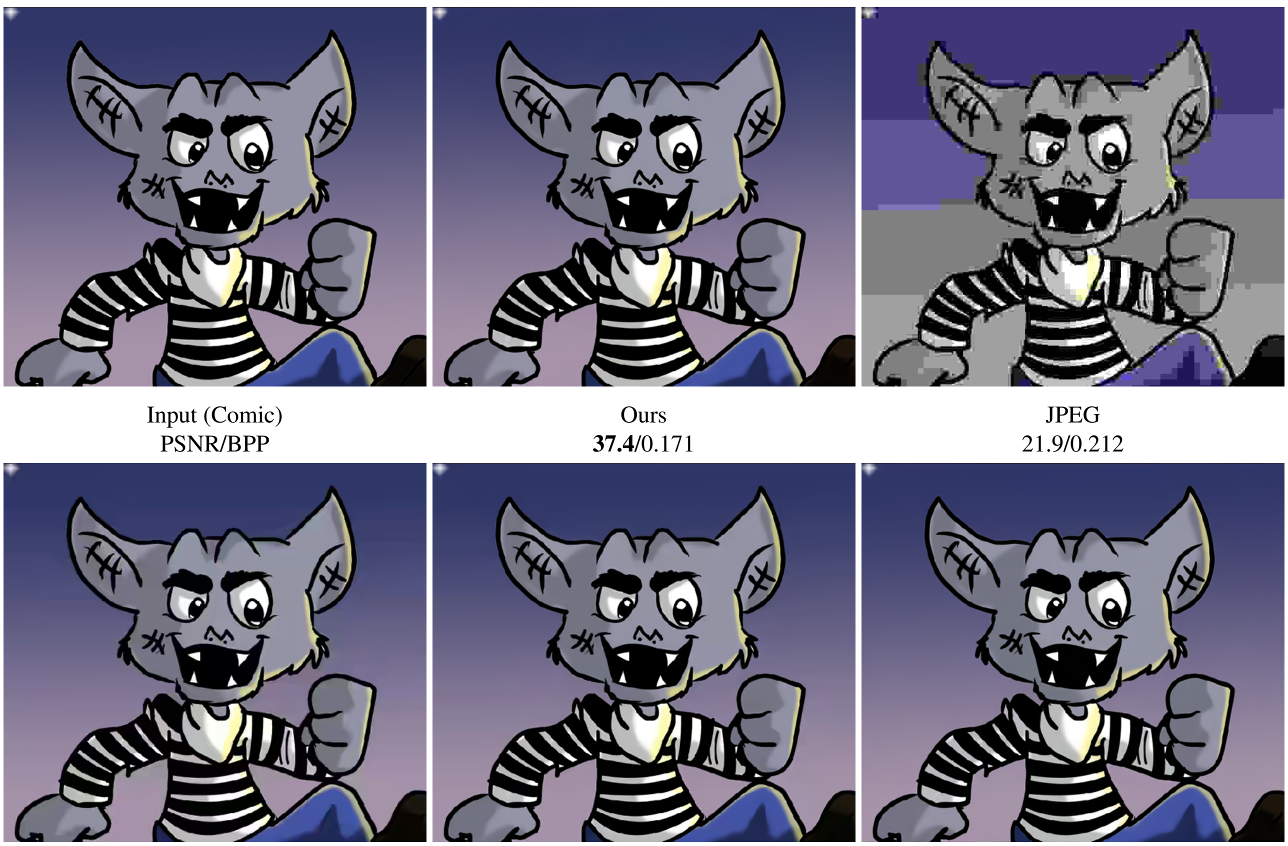

Input (Comic) PSNR/BPP | Ours **37.4**/0.171 | JPEG 21.9/0.212

BPG 32.6/0.177 | VVC 37.2/0.183 | WACNN [5] 36.7/0.180

Figure 2: Qualitative comparison with other compression methods on a comic image.

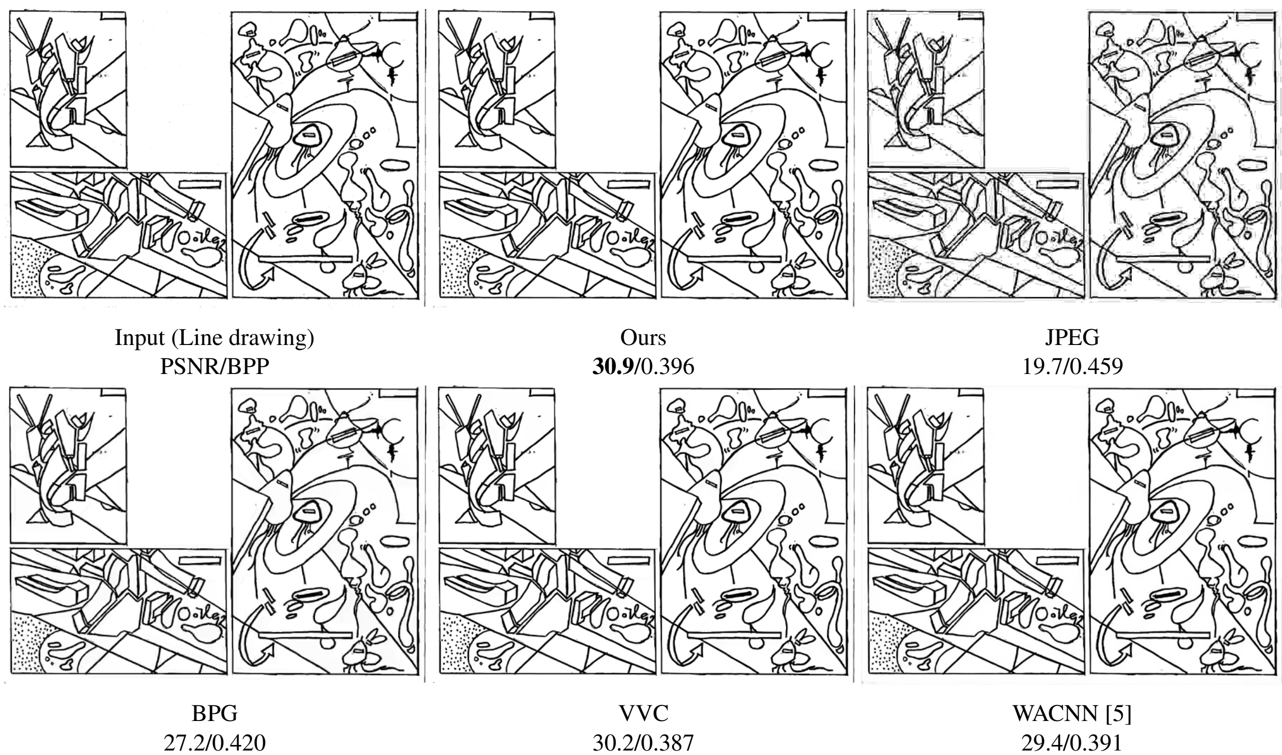

Figure 3: Qualitative comparison with other compression methods on a line drawing.

Figure 4: Qualitative comparison with other compression methods on a vector art.

Input	Ours	Without adapters
PSNR/BPP	35.4/0.099	34.5/0.090

Figure 5: Qualitative results on the effectiveness of adapters on a vector art.